Editor: **Daniel B. Neill,** H.J. Heinz III College, Carnegie Mellon University, neill@cs.cmu.edu

How Social Media Will Change Public Health

Mark Dredze, *Johns Hopkins University*

Social media such as Twitter have created platforms for people to broadcast information, thoughts, and feelings about their daily lives. Since Twitter messages (called *tweets*) often reflect in-the-moment updates, they're filled with useful observations and information about the larger world. Researchers have examined a range of applications based on tweets, ranging from political polling[1] to earthquake monitoring,[2] that have demonstrated Twitter's ability to deliver fast, cheap, and reliable tools for monitoring real-world events.

These successes have drawn interest from the public-health community, whose goal is to study the health of a population and develop policies that improve health outcomes. Traditionally, this requires expensive, time-consuming monitoring mechanisms, primarily surveys and data collection from clinical encounters. Even high-priority projects, such as the US Centers for Disease Control and Prevention's (CDC's) FluView program that tracks the weekly US influenza rate, are still slow because they require clinical data aggregation. Twitter and other social media could reduce cost and provide real-time statistics about public health.

Recent work in machine learning and natural language processing has studied the health content of tweets and demonstrated the potential for extracting useful public-health information from their aggregation. This article examines the types of health topics discussed on Twitter, and how tweets can both augment existing public-health capabilities and enable new ones. I also discuss key challenges that researchers must address to deliver high-quality tools to the public-health community.

Discovering Health Topics on Twitter

Twitter's size and breadth make it difficult to determine exactly which types of public-health work it can support. Initial work in my research group[3,4] explored health-related tweets and topics on Twitter through the development of new computational models. Because many public-health activities are disease-oriented, we developed a model that discovered diseases (ailments) from raw tweets for guided exploration, rather than relying on predefined illnesses. We used supervised learning to filter tweets and find health-related messages, yielding 1.6 million English health tweets from March 2009 to October 2010.

To explore these tweets, we developed the Ailment Topic Aspect Model (ATAM), a probabilistic graphical model for uncovering ailments.[3] ATAM assumes that each message discusses a single ailment, manifested through the message's words, and associates three types of words (general disease words, symptoms, and treatments) with ailments. For example, the message "fever + headache = flu, home sick with Tylenol" discusses influenza, where "fever" and "headache" are symptoms, "Tylenol" a treatment, and "flu" a general word associated with the ailment.

Human annotators labeled 15 ailments discovered by ATAM, including headaches, influenza, insomnia, obesity, dental problems, and seasonal allergies. Examining the words, symptoms, and treatments most associated with each ailment, and the groups of messages that discuss each ailment, can support a variety of public-health initiatives.

Augmenting Existing Public-Health Capabilities

A core capability of public-health programs, biosurveillance monitors a population for adverse health events, which include expected seasonal events, such as influenza or environmental allergies, disease outbreaks, such as the H1N1 virus, and other health threats, such as food poisoning or a biochemical contaminant. Surveillance is the

1541-1672/12/$31.00 © 2012 IEEE
Published by the IEEE Computer Society

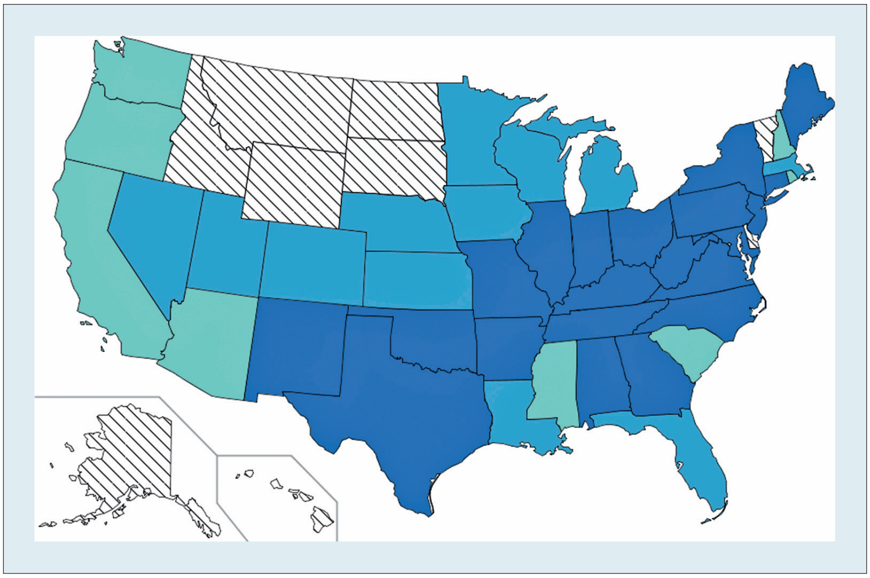

Figure 1. The rate of Twitter messages about seasonal allergies for June 2010. Messages were automatically coded using a machine-learning method and geo-located based on user-provided location. Overall shading indicates significant allergy messages, showing the heart of allergy season. States in the Northeast and Midwest are particularly active. Dashed states had insufficient data.

key first step in any comprehensive response strategy. Consider the example of the H1N1 virus, which struck the US in 2009. Public-health officials had to direct vaccine supplies to the areas and populations where they were most needed, requiring accurate information about where H1N1 infections were occurring and which demographic groups were most affected.

Traditional biosurveillance relies on information collected from clinical encounters, a time-consuming process. For example, in the case of influenza tracking, the CDC requires two weeks to collect and release statistics about the US flu rate. Web-based approaches can produce faster results, such as Google Flu Trends,[5] which analyzes real-time search queries to produce a daily flu rate. When users search for flu-related queries, such as "flu medicine" or "flu symptoms," Google aggregates these statistics to measure rises in flu traffic. These have been shown to correlate with the CDC's official estimates, providing more timely influenza surveillance.[5]

Analyzing Twitter messages could provide a similar surveillance capability. Studies[4,6] have shown correlations between influenza tweets and CDC data, using supervised learning and unsupervised learning. This idea has been extended to low-resource settings in developing countries, such as surveillance of cholera in Haiti.[7]

Because Twitter provides location information for some tweets, biosurveillance can be geographically localized. For example, we visualized the per capita tweeting rate about seasonal allergies for the month of June 2010 (in Figure 1, where the darker colors indicate more tweets[4]). As expected, the Midwest and Northeast have substantial Twitter traffic as compared to other regions of the US, which follows the expected start of allergy season. By contrast, the winter months have few allergy messages.

Beyond surveillance, Twitter can support other public-health tasks, such as health risk assessments. For example, the annual CDC Behavioral Risk Factors Study surveys more than 300,000 people nationwide for several risk factors, such as asthma, smoking, and exercise. The study is both expensive and time-consuming, making it inappropriate for rapid hypothesis generation and testing. Twitter could augment this survey by investigating additional questions or providing faster results. We compared[4] each of the survey questions that had corresponding ailments discovered by ATAM across the 50 states, uncovering interesting correlations, such as a positive correlation between states with high smoking rates and those with high Twitter message rates about cancer (r = 0.648), a negative correlation between exercise and obesity messages (r = −0.201), and a negative correlation between good healthcare coverage and messages about ailments in general (r = −0.253).

Creating New Public-Health Capabilities

The monitoring of Twitter data can also enable the creation of entirely new public-health capabilities, supported by both the expressiveness of tweets and the coverage of topics not normally included in public-health data, particularly those that people are reluctant to discuss with healthcare workers.

The public forum of social media encourages messages that express a range of details, yielding health information such as the illness, symptoms, and treatment strategy—for example, "took some Tylenol for my flu" or "stuck home with flu and 102 fever." Consider Figure 2, which shows the word cloud for insomnia generated via ATAM output, in which word size corresponds to influenza likelihood, and color indicates word type (red are symptoms, green are treatments, and blue are general words). Although search engine users might

turn to Google to look for insomnia remedies, Twitter users provide a variety of details about their sleepless nights.

Additionally, Twitter covers different topics than those covered by traditional public-health data sources such as clinical encounters and phone surveys. Behaviors that people might be reluctant to share with physicians are on full display on Twitter, including behaviors, opinions, and subpopulations that are otherwise difficult to track through traditional mechanisms, suggesting a whole new area of large-scale public-health research.

Disease self-management can be hard to study, as it doesn't involve a physician and patients might be reluctant to share unapproved practices with health officials. We studied medication usage from tweets by creating medication usage profiles based on ailment groupings.[4] For pain relievers, for example, we found that Tylenol and Advil have broad profiles (headache, cold relief, and so on) while Vicodin is targeted at dental problems and injuries. For allergy medication, Claritin and Zyrtec were almost exclusively used to treat allergies, while off-label uses of Benadryl included insomnia. Monitoring medication usage on Twitter can discover new trends in self-medication otherwise unreported by patients.

The information gap in traditional public health is especially prevalent in patient-directed programs such as weight loss and smoking cessation. These depend on a sustained effort from patients outside the clinical setting, making it difficult to track and measure patient efforts. For example, a recent study of 15,000 tweets found that Twitter is commonly used to manage and share information about health-promoting physical

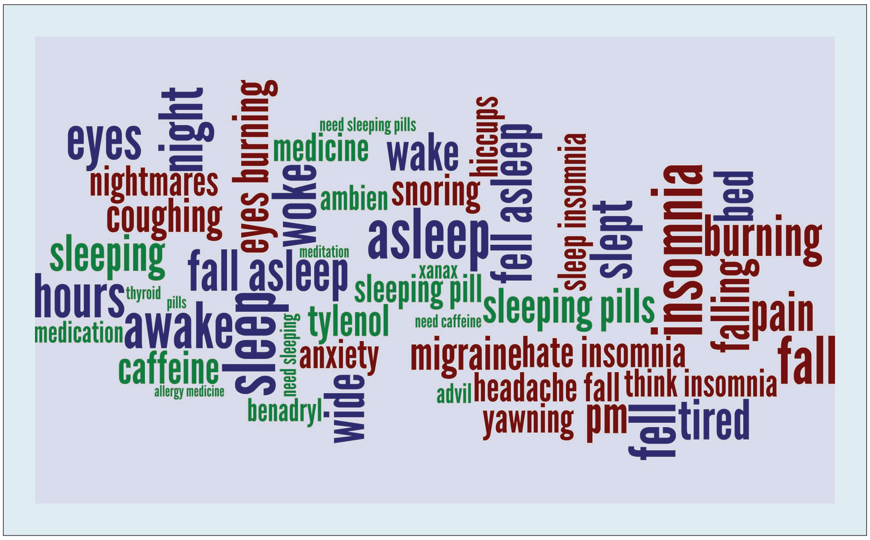

Figure 2. A word cloud visualization showing the words most associated with the ailment "insomnia" as discovered by a machine-learning model that examined 1.6 million tweets related to health. Larger fonts indicate more related terms, blue indicates general terms, red highlights symptoms, and green represents treatments. General words such as "hours," "awake," and "tired" characterize insomnia messages, with symptoms such as "nightmares" and "yawning" and treatments of "Benadryl" and "sleeping pills."

activities.[8] Tweets focused on exercise included muscle-strengthening, aerobic, and flexibility-enhancing activities. An analysis of the content revealed that most tweets reported evidence of or plans for exercising. The frequency of such messages suggests that the social supports provided by Twitter could be used as a platform for encouraging exercise and health-promoting behavior. Additionally, roughly 10 percent of the messages posed exercise questions to other users, and many contained advertisements for a product or service. Mining this information could reveal trends of physical activity, as well as new ways of promoting healthy behaviors.

Although dental pain is a common problem, only a few complaints result in seeing a dentist, and thus clinical evidence covers only a small part of the population. A recent study considered reports of dental pain on Twitter as a means of surveying a larger spectrum of patients.[9] A survey of 772 messages revealed a variety of topics, including reporting dental pain, action taken in response, impact on daily life, and advice sought from the Twitter community. More than 80 percent of messages discussed general pain, 22 percent discussed taking some responsive action, and 15 percent discussed impact on daily activities. While actions included seeing a dentist (44 percent), just as many self-medicated (43 percent). These findings show that Twitter can broaden the study population and indicate that effective treatment of dental problems relies on providing accurate information about self-management. The prevalence of dental communications suggests that Twitter can be an effective medium for dental professionals to disseminate self-management information.

Twitter is of special interest in studies of patient health behaviors, such as drug and alcohol use. Kyle Prier and colleagues explored tweets about smoking to learn more about tobacco use.[10] They used an unsupervised clustering algorithm to group smoking tweets, discovering themes that

reflected general substance abuse, addiction recovery, tobacco promotion (bars and clubs), and antismoking advertising campaigns. These themes suggest Twitter as a promising data source for tobacco use behaviors and trends.

Each of these studies—on self-medication, exercise, dental pain, and smoking—demonstrates the potential for new areas of public-health research based on Twitter data. The next few years promise detailed clinical studies using data in each of these areas as well as whole new types of questions. How will Twitter data impact the study of community health behaviors, mental health, and biosurveillance customized to specific demographic groups? The development of new technologies coupled with the exploration of these questions has great potential to expand the scope of public-health practice.

Automatically extracting meaning from text with computational tools is difficult, particularly when the text is multilingual and informal. Yet computational algorithms will be required for practical public health applications of Twitter data. Preliminary research suggests that aggregating millions of messages can resolve difficulties in understanding individual messages: the tweet "headache" is ambiguous, but a corresponding increase in messages of the form "bad headache, home sick with flu" suggests a common cause. However, a deeper analysis of individual tweets, which might be required for some tasks, remains a challenging problem.

In addition, bias pervades all stages of social media analysis: social media users aren't representative of the entire population, user groups may be more or less inclined to tweet information, and users might report inaccurate or imprecise diagnoses (for example, influenza instead of H1N1). Controlling for bias is a hallmark of clinical research, yet social media biases are little understood. Large-scale aggregation could help obviate biases for common illnesses. However, smaller populations will require bias correction, which may rely on automatically inferring user demographics for sampling adjustments.

Finally, social media research must consider user privacy. Even when data are publicly available, users have privacy expectations, such as concern over algorithms that infer unstated user demographics or diagnoses from public data. Although studies have posed little concern so far, an increase in research complexity could cause a commensurate rise in legal and ethical issues. Social media researchers must remain vigilant regarding privacy issues.

Regardless, with the development of new technologies addressing these challenges, we can expect to see entirely new capabilities for public-health research, policy, and practice.

References

1. B. O'Connor et al., "From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series," *Proc. Int'l Conf. Weblogs and Social Media* (ICWSM), AAAI, 2010, pp. 122–129.
2. T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake Shakes Twitter Users: Real-Time Event Detection by Social Sensors," *Proc. World Wide Web Conf.* (WWW 10), ACM, 2010, pp. 851–860.
3. M.J. Paul and M. Dredze, "A Model for Mining Public Health Topics from Twitter," tech. report, Johns Hopkins Univ., 2011.
4. M.J. Paul and M. Dredze, "You Are What You Tweet: Analyzing Twitter for Public Health," *Proc. Int'l Conf. Weblogs and Social Media* (ICWSM), AAAI, 2011, pp. 265–272.
5. J. Ginsberg et al., "Detecting Influenza Epidemics Using Search Engine Query Data," *Nature*, vol. 457, no. 7232, 2008, pp. 1012–1014.
6. A. Culotta, "Lightweight Methods to Estimate Influenza Rates and Alcohol Sales Volume from Twitter Messages," *Language Resources and Evaluation*, 2012; www.springerlink.com/content/x738116375268725/.
7. R. Chunara, J.R. Andrews, and J.S. Brownstein, "Social and News Media Enable Estimation of Epidemiological Patterns Early in the 2010 Haitian Cholera Outbreak," *Am. J. Tropical Medicine and Hygiene*, vol. 86, no. 1, 2012, pp. 39–45.
8. L. Kendall et al., "Descriptive Analysis of Physical Activity Conversations on Twitter," *Proc. Extended Abstracts on Human Factors in Computing Systems (CHI),* ACM, 2011, pp. 1555–1560.
9. N. Heaivilin et al., "Public Health Surveillance of Dental Pain via Twitter," *J. Dental Research*, vol. 90, no. 9, 2011, pp. 1047–1051.
10. K.W. Prier et al., "Identifying Health-Related Topics on Twitter: An Exploration of Tobacco-Related Tweets as a Test Topic," *Social Computing, Behavioral-Cultural Modeling, and Prediction*, LNCS 6589, Springer, 2011, pp. 18–25.

Mark Dredze is an assistant research professor in computer science and a research scientist at the Human Language Technology Center of Excellence at Johns Hopkins University. Dredze has a PhD in computer science from the University of Pennsylvania. Contact him at mdredze@cs.jhu.edu.

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*